

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Barillaud et al.	§	Group Art Unit: 2445
	§	
Serial No. 09/714,724	§	Examiner: Swearingen, Jeffrey R.
	§	
Filed: November 16, 2000	§	Confirmation No.: 9479
	§	
For: Method and System for Automatic	§	
Load Balancing of Advertised Services	§	
by Service Information Propagation	§	
Based on User On-Demand Requests		

35525

PATENT TRADEMARK OFFICE
CUSTOMER NUMBER

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF (37 C.F.R. 41.37)

This brief is in furtherance of the Notice of Reinstatement of Appeal, filed in this case on March 10, 2010.

No fee is believed to be required for filing this Appeal Brief, as the appeal has been reinstated pursuant to MPEP 1204.01. No additional fees are believed to be required. If, however, any fees are required, I authorize the Commissioner to charge these fees which may be required to IBM Corporation Deposit Account No. 09-0447. No extension of time is believed to be necessary. If, however, an extension of time is required, the extension is requested, and I authorize the Commissioner to charge any fees for this extension to IBM Corporation Deposit Account No. 09-0447.

REAL PARTY IN INTEREST

The real party in interest in this appeal is the following party: International Business Machines Corporation of Armonk, New York.

RELATED APPEALS AND INTERFERENCES

This appeal has no related proceedings or interferences.

STATUS OF CLAIMS

A. TOTAL NUMBER OF CLAIMS IN APPLICATION

The claims in the application are: 1-41

B. STATUS OF ALL THE CLAIMS IN APPLICATION

Claims canceled: 1-2 and 13-36

Claims withdrawn from consideration but not canceled: None

Claims pending: 3-12 and 37-41

Claims allowed: None

Claims rejected: 3-12 and 37-41

Claims objected to: None

C. CLAIMS ON APPEAL

The claims on appeal are: 3-12 and 37-41

STATUS OF AMENDMENTS

This appeal is a reinstated appeal being filed pursuant to MPEP 1204.01, and no amendments have been filed in response to the Final Office Action dated January 6, 2010. Amendments have been filed in response to the Office Action dated June 2, 2009, which re-opened prosecution of this case when it was under a previous appeal pursuant to an Appeal Brief filed by Appellants on February 18, 2009.

SUMMARY OF CLAIMED SUBJECT MATTER

In a highly distributed computational system, the applications that perform operations for a given network service may be dispersed on physical devices throughout the network. Applications on other physical devices that desire access to the given network service must be provided with information on the manner in which a connection to the network service can be obtained. A complete inventory of available networked resources may be distributed throughout the system.

In any given network, the demand for networked resources fluctuates over time. Generally, network management software within the distributed computing system satisfies the demand for network resources using some type of load balancing such that all service requesters eventually get access to the requested service. It is sometimes critical to load balance the demand for services by distributing the request workload across the entire system in order to ensure fair access. The pending claims provide access to a *plurality* of different network *services* in a heterogeneous distributed environment, where demand balancing for network services is achieved architecturally, by having “each client uniquely associated with a *local* service manager”, and “each *local* service manager uniquely associated with a *distributed* service manager” to select among available services.

A. CLAIM 3 - INDEPENDENT

The subject matter of Claim 3 is directed to a method of balancing demand for networked services in a distributed data processing system (Specification page 11, lines 1-3; Figure 1, element 110). One or more local service managers within the distributed data processing system are initialized, where each local service manager has information about, and provides access to, networked services defined within a respective local region of the distributed data processing system for clients within the distributed data processing system (Specification page 4, lines 8-11; page 12, lines 12-17; Figure 3, element 322), where each client is uniquely associated with a local service manager (Specification page 4, lines 7-8; Figure 5, elements 561/563; Figure 7, elements 721/723; Figure 8, elements 821/823; Figure 9A-9B, elements 921/923; Figure 10A-10B, elements 1021/1023, 1031/1033 and 1051/1053; Figure 12, elements 1241/1243). In addition, one or more distributed service managers within the distributed data processing system are initialized, where each distributed service manager provides access to the networked services to the local service managers within the distributed data processing system (Specification page 4,

lines 13-17; page 12, lines 17-23; Figure 3, element 324), where each local service manager is uniquely associated with a distributed service manager (Specification page 4, lines 11-13; page 13, lines 27-28; Figure 6). A request for a networked service - from a local service manager for which the local service manager lacks information – is received at a distributed service manager (Specification page 4, lines 17-24; page 16, lines 27-30; Figure 4C, element 440). A determination is made as to whether the distributed service manager has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service, by referencing a cache maintained by the distributed service manager which contains information resulting from prior requests for networked services (Specification page 16, line 30 – page 17, line 4; Figure 4C, element 442). Information for referencing a matched networked service is returned (Specification page 18, lines 24-28; Figure 4C, element 456).

B. CLAIM 37 – INDEPENDENT

The subject matter of Claim 37 is directed to a method of balancing demand for networked services in a distributed data processing system (Specification page 11, lines 1-3; Figure 1, element 110). The method comprises the step of initializing one or more local service managers within the distributed data processing system, wherein each local service manager has information about and provides access to networked services defined within a respective local region of the distributed data processing system for clients within the distributed data processing system (Specification page 4, lines 8-11; page 12, lines 12-17; Figure 3, element 322), and wherein each client is uniquely associated with a local service manager (Specification page 4, lines 7-8; Figure 5, elements 561/563; Figure 7, elements 721/723; Figure 8, elements 821/823; Figure 9A-9B, elements 921/923; Figure 10A-10B, elements 1021/1023, 1031/1033 and 1051/1053; Figure 12, elements 1241/1243). The method comprises the step of initializing one or more distributed service managers within the distributed data processing system, wherein each distributed service manager provides access to the networked services to the local service managers within the distributed data processing system (Specification page 4, lines 13-17; page 12, lines 17-23; Figure 3, element 324), and wherein each local service manager is uniquely associated with a distributed service manager (Specification page 4, lines 11-13; page 13, lines 27-28; Figure 6). The method comprises the step of receiving, at a distributed service manager, a

request for a networked service from a local service manager for which the local service manager lacks information (Specification page 4, lines 17-24; page 16, lines 27-30; Figure 4C, element 440). The method comprises the step of determining whether the distributed service manager has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service, wherein the determining step is accomplished by reference to a cache maintained by the distributed service manager which contains information resulting from prior requests for networked services (Specification page 16, line 30 – page 17, line 4; Figure 4C, element 442). The method comprises the step of returning information for referencing a matched networked service (Specification page 18, lines 24-28; Figure 4C, element 456). The method comprises the step of configuring the local service manager to not provide access to object request broker (ORB) services that provide internal service and which are valid only in a scope of a local ORB (Specification page 11, lines 25-27; page 15, lines 2-3). The method comprises the step of configuring the local service manager to provide access to ORB services that are instantiated on each ORB only through requests based on an ORB identifier (Specification page 11, lines 28-29; page 15, lines 3-5). The method comprises the step of configuring the local service manager to provide access to ORB services that may be accessed from outside the scope of the local ORB through requests based on both a service specification string and an ORB identifier (Specification page 11, lines 29-31; page 15, lines 5-7).

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The grounds of rejection to review on appeal are as follows:

A. GROUND OF REJECTION 1

The rejection of Claim 3 under 35 U.S.C. § 102(b) as being anticipated by Derby et al. (U.S. Patent No. 5,426,637);

B. GROUND OF REJECTION 2

The rejection of Claims 3-7, 38 and 39 under 35 U.S.C. § 103 as being unpatentable over Elnozahy et al. (U.S. Patent No. 6,014,686) in view of Chandra et al. (U.S. Patent No. 6,457,047);

C. GROUND OF REJECTION 3

The rejection of Claims 8-12 and 41 under 35 U.S.C. § 103 as being unpatentable over Elnozahy et al. (U.S. Patent No. 6,014,686) in view of Chandra et al. (U.S. Patent No. 6,457,047) as applied to Claim 8, and further in view of Jindal et al. (U.S. Patent No. 6,324,580); and

D. GROUND OF REJECTION 4

The rejection of Claim 37 under 35 U.S.C. § 103 as being unpatentable over Elnozahy et al. (U.S. Patent No. 6,014,686) in view of Chandra et al. (U.S. Patent No. 6,457,047) as applied to Claim 3, and further in view of Fowlow et al. (U.S. Patent No. 5,920,868).

ARGUMENT

A. GROUND OF REJECTION 1 (Claim 3)

Claim 3 stands rejected under 35 U.S.C. § 102(b) as being anticipated by Derby et al. (U.S. Patent No. 5,426,637), hereinafter “Derby”.

To anticipate, the prior art must teach all the claim elements and the claimed arrangement. Section 102 embodies the concept of novelty—if a device or process has been previously invented (and disclosed to the public), then it is not new, and therefore the claimed invention is “anticipated” by the prior invention. . . . Because the hallmark of anticipation is prior invention, the prior art reference—in order to anticipate under 35 U.S.C. § 102—must not only disclose all elements of the claim within the four corners of the document, but must also disclose those elements “arranged as in the claim.” Focusing for a moment on *arrangement* – to anticipate, the reference must teach “all of the limitations arranged or combined in the same way as recited in the claim.” *Net MoneyIn v. Verisign*, 545 F.3d 1359 (Fed. Cir. 2008).

Generally, Claim 3 is directed to a method of *balancing demand for network services*. The cited Derby reference is instead directed to a technique for interconnecting LANs to a WAN, and provides no type of *network service demand balancing*, as per Claim 3.

With respect to Claim 3, such claim recites “initializing one or more local service managers within the distributed data processing system, *wherein each local service manager has information about and provides access to networked services* defined within a respective local region of the distributed data processing system for clients within the distributed data processing system, and *wherein each client is uniquely associated with a local service manager*” (emphasis added). As can be seen, each local service manager (i) has information about and (ii) provides access to *networked services*. In addition, *each client is uniquely associated with a local service manager*.

In rejecting Claim 3, the Examiner alleges that Derby teaches the claimed local service manager at col. 6, line 10. Appellants show that there, Derby describes a ‘LAN access agent’. However, this Derby ‘LAN access agent’ is not equivalent to the claimed ‘local service manager’ since the Derby ‘LAN access agent’ is not described as having a client uniquely associated with it. Therefore, Derby does not teach “*wherein each client is uniquely associated with a local service manager*” as a part of a network services demand balancing technique, and therefore

Derby does not anticipate Claim 3 as such reference does not teach all of the limitations arranged or combined in the same way as recited in the claim.

This failure of the cited reference to teach such ‘uniquely associated’ feature can also be seen at Derby Figure 10, where multiple WAN agents (that include the ‘access agent’ that is alleged as being equivalent to the claimed ‘local service manager’) are associated with a given LAN, as described at col. 14, lines 30-32. The WAN access node does not interact with individual clients – but instead with the LAN - and thus there would be no need for the Derby LAN access agent (which is a part of the WAN node as shown in Figure 2) to have a client uniquely associated with it since the function of this WAN is to interconnect the LANs together (Derby col. 2, line 63 – col. 3, line 16). Therefore, it is further shown that the Derby ‘LAN access agent’ is not described as having a client uniquely associated with it.

Further with respect to Claim 3, and as alluded to above, the claimed local service manager *has information about and provides access to networked services defined within a respective local region* of the distributed data processing system. The cited reference does not describe that its LAN access agent – which the Examiner alleges as being equivalent to the claimed local service manager – *has information about and provides access to networked services defined within a respective local region* for at least two reasons. First, the Derby LAN access agent is not described as having any information about networked services. Secondly, the Derby LAN access agent is not described as providing access to networked services *defined within a respective local region*. In fact, Derby is keen on merely providing a conduit for data traveling long distances between a plurality of LANs and a WAN (Derby col. 5, lines 30-53), and therefore would have no need to provide access to *local* region network services due to the desire to provide *remote* access for all device requests. Thus, it is further urged that Claim 3 has been erroneously rejected as Derby does not teach all of the limitations arranged or combined in the same way as recited in the claim.

In an apparent acknowledgement of this missing claimed characteristic with respect to the claimed ‘local service manager’, the Examiner now newly provides a new interpretation of Derby in the “Response to Arguments” section on page 2 of the Final Office Action dated January 6, 2010. The Examiner now newly states that the Derby ‘directory services unit 22’ “*has information about and provides access to networked services defined within a respective local region*” aspect of Claim 3. It is initially noted that this directory services unit 22 is not a part of

the Derby LAN access agent 16. It is further noted that the directory services unit is described as containing information about ‘reachability’ of *stations* (col. 6, lines 26-32), and not information pertaining to *networked services defined within a respective local region*, as claimed – further evidencing that Claim 3 has been erroneously rejected. The Examiner further opines that since the directory services unit receives necessary addressing information from the local LAN access agent, the local LAN access agent must possess information about networked services (see top of page 3 of the Final Office Action dated January 6, 2010). This implicit ‘inherency’ assertion is also not true for several reasons. First, only physical device addressing information is described by the Derby directory services unit, but no underlying service information, as previously shown. Second, an address does not imply a service. The fact that the Derby LAN access agent provides an address does not mean it has knowledge of any service, as such knowledge would more likely be at a requesting client (to the extent the client is requesting a service). The LAN information received by the LAN access agent is received from the LAN itself – and not the client – and therefore such received information (by the LAN access agent) does not inherently include information about services. Therefore, this new interpretation of Derby also fails to establish that Derby teaches all of the limitations arranged or combined in the same way as recited in the claim.

Further with respect to Claim 3, such claim recites, “receiving, at a distributed service manager, a request for a networked service from a local service manager for which the local service manager lacks information”. As can be seen, a request for a networked service is received at a distributed service manager, where such request is received from a local service manager for which *the local service manager lacks information*.

In rejecting this aspect of Claim 3, the Examiner cites Derby’s teaching at col. 7, lines 24-40 as teaching all aspects of such networked service request. Appellants show that there, Derby states:

For example, when the LAN protocol initiates a search procedure, the protocol agent 18 registers, if present, the reachability information for the local resources embedded in the respective protocol frames (i.e., the address prefixes) with directory services unit 22 and initiates **searches for resources** which are not available locally. Reciprocally, directory services unit 22 invokes the protocol components 18 in the appropriate access nodes in order to search the local LANs for the **location of a particular destination LAN station**. The results of such searches may then be cached by the WAN access agents in address cache 24 connected to protocol component 18, used to expedite the processing of future search procedures. That is, address cache 24 stores network layer addresses of

network layer entities residing in stations attached to local and remote LANs, along with supporting information.

As can be seen, this cited passage describes a search procedure that searches for hardware resources (“LAN station”). This cited passage does not describe any type of request for a networked ‘service’. This can also be seen by Derby col. 12, lines 9-11 (“search request for the identified **destination**”); if a destination match is found, *the MAC address of the device is returned* (col. 12, lines 35-37); the information maintained in connection tables is specially directed to connections between hardware resources (“LAN stations”, per col. 7, lines 33-34 and lines 34-48) - and not to any type of ‘services’ that may be available to such hardware resources. Thus, this cited Derby passage does not describe, “receiving, at a distributed service manager, *a request for a networked service from a local service manager for which the local service manager lacks information*”, as claimed.

Further with respect to Claim 3, such claim recites, “determining whether the distributed service manager has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service, wherein the determining step is accomplished by reference to a cache maintained by the distributed service manager which contains information resulting from prior requests for networked services”. As can be seen, a determination is made whether the distributed service manager – which the Examiner alleges is equivalent to Derby’s ‘directory services unit’ – *has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service*.

The Examiner alleges that Derby teaches all such distributed service manager determination aspects – including a determining being made as to whether such distributed service manager has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service – at col. 7, lines 24-40. Appellants show that there, Derby states:

For example, when the LAN protocol initiates a search procedure, the protocol agent 18 registers, if present, the reachability information for the local resources embedded in the respective protocol frames (i.e., the address prefixes) with directory services unit 22 and initiates searches for resources which are not available locally. Reciprocally, **directory services unit 22 invokes the protocol components 18 in the appropriate access nodes**

in order to search the local LANs for the location of a particular destination LAN station. The results of such searches may then be cached by the WAN access agents in address cache 24 connected to protocol component 18, used to expedite the processing of future search procedures. That is, address cache 24 stores network layer addresses of network layer entities residing in stations attached to local and remote LANs, along with supporting information.

As can be seen, this cited passage does not describe any determination being made as to whether the Derby ‘directory services unit’ – which the Examiner alleges is equivalent to the claimed ‘distributed service manager’ - *has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service.* Instead, this cited passage merely states that this ‘directory services unit’ performs a searching operation with respect to remote nodes. No determination is made as to whether the directory services unit itself *has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service.* Thus, it is further urged that Claim 3 is not anticipated by Derby, and thus Claim 3 has been erroneously rejected.

This failure to teach all of the limitations is even expressly acknowledged by the Examiner themselves. For example, on page 2 of the Final Office Action dated January 6, 2010, the Examiner states that “It is *commonly known* to one of ordinary skill in the art that while a client must be singularly associated with a server that is an access point for a network, the server may function as an access point for multiple clients to the network” – which is an ‘obviousness’ rationale and yet Claim 3 is being rejected under 35 USC 102. Such statement is effectively an admission that the cited reference does not teach this ‘unique’ association due to the Examiner’s reliance on what is commonly known to those of ordinary skill in the art instead of what the anticipatory reference actually teaches. *In addition*, this Derby reference does not teach that each ‘client’ is ‘*uniquely* associated’ with its ‘LAN access agent’ (that the Examiner alleges is equivalent to the claimed ‘local service manager’). Instead, this ‘LAN access agent’ is described as forming a point of attachment between the WAN and the *attached external LANs* in order to mediate between the external LAN protocols and the protocols available on the backbone WAN (col. 6, lines 50-55). This cited reference *also* does not describe that each ‘LAN access agent’ – which the Examiner alleges is equivalent to the claimed ‘local service manager’ – is ‘*uniquely* associated with’ the ‘directory services’ (which the Examiner alleges is equivalent to the claimed ‘distributed service manager’, as previously shown in detail and as graphically depicted by

Derby's Figure 10).

B. GROUND OF REJECTION 2 (Claims 3-7, 38 and 39)

Claims 3-7, 38 and 39 stand rejected under 35 U.S.C. § 103 as being unpatentable over Elnozahy et al. (U.S. Patent No. 6,014,686), hereinafter "Elnozahy" in view of Chandra et al. (U.S. Patent No. 6,457,047), hereinafter "Chandra".

The Examiner bears the burden of establishing a *prima facie* case of obviousness based on prior art when rejecting claims under 35 U.S.C. § 103. *In re Fritch*, 972 F.2d 1260, 23 U.S.P.Q.2d 1780 (Fed. Cir. 1992). *In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). Only if that burden is met, does the burden of coming forward with evidence or argument shift to the applicant. *Id.* All words in a claim must be considered in judging the patentability of that claim against the prior art." MPEP 2143.03; *In re Wilson*, 424 F.2d 1382, 1385, 165 USPQ 494, 496 (CCPA 1970). If the examiner fails to establish a *prima facie* case, the rejection is improper and will be overturned. *In re Fine*, 837 F.2d 1071, 1074, 5 USPQ2d 1596, 1598 (Fed. Cir. 1988). In the absence of a proper *prima facie* case of obviousness, an applicant who complies with the other statutory requirements is entitled to a patent. *See In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992).

1. Claims 3, 5 and 6

With respect to Claim 3, such claim recites both a local service manager (one or more) and a distributed service manager (one or more), where each distributed service manager provides access to networked services to the local service manager(s). Importantly, *each of these distributed service managers provides access to networked services to the local service managers.*

In rejecting Claim 3, the Examiner alleges that Elnozahy teaches both types of managers by the 'Cell Directory Service' and the 'CDS server'. Appellants urge that *these two cited items are the **same thing***, as the 'CDS' portion of 'CDS server' stands for 'Cell Directory Service', as described by Elnozahy at col. 1, line 49. Importantly, Elnozahy does not describe a Cell Directory Service that is separate from the CDS server – instead the CDS server is what is used in providing the Elnozahy Cell Directory Service, as described by Elnozahy at col. 2, line 56 – col.

3, line 13 and col. 5, lines 23 (“The highly available **CDS** in accordance with the principles of the present invention **consist of two main components: standard CDS servers** 62 and 72 and high availability agents 68 and 76”).

As yet another example of how the Examiner is impermissibly double counting the described Cell Directory Service that is provided by the CDS server – where the Examiner alleges that the Cell Directory Service (CDS) server reads on both the ‘local service manager’ as well as the ‘distributed service manager’ - the cited Elnozahy reference does not describe two separate initialization steps for each of its ‘CDS server’ and ‘CDS’ as they are one and the same thing. In contrast, Claim 3 recites two separate initialization steps, with one pertaining to local service manager(s) and another pertaining to distributed service managers.

As yet another example of how the ‘Elnozahy CDS’ server of ‘Cell Directory Service’ does not teach/suggest both the claimed ‘local service manager’ as well as the ‘distributed service manager’, such claim recites “wherein **each distributed service manager provides access to the networked services to the local service managers** within the distributed data processing system, and wherein each local service manager is uniquely associated with a distributed service manager” (emphasis added). The cited Elnozahy reference does not describe that its ‘CDS server’ (which is alleged as being equivalent to the claimed ‘distributed service manager’) provides access to networked services to a ‘Cell Directory Service’ (which is also alleged as being equivalent to the claimed ‘local service manager’).

In the Examiner’s rebuttal remarks to this these identified prima facie obviousness deficiencies, the Examiner notes that the cited reference does in fact teach two different CDS servers since it teaches a ‘master’ CDS server and a ‘replica’ CDS server (see the top of page 4 of the Final Office Action dated January 6, 2010). Appellants urge that these master and replica CDS servers are not hierarchically configured per the claimed distributed service manger and local service manager, since one of these CDS servers does not provide access to the networked servers *to the other* of these CDS servers. Per Claim 3, “wherein **each distributed service manager provides access to the networked services to the local service managers** within the distributed data processing system. Indeed, per the teachings of the cited reference, the redundant CDS servers provide a server ‘backup’ in the event of a primary server failure (Elnozahy col. 3, lines 42-44). These master and replica servers do not provide the hierarchy

recited in Claim 1, whereby (i) a *local service manager* provides access to networked services with a local region *for clients*, and (ii) a *distribute service manager* provides access to networked services to such a local service manager – which is the ‘hierarchy’ aspect of Claim 3 that is not taught/suggested by the cited references. Indeed, the master CDS server is ***not aware of the other servers*** – so such hierarchical configuration between multiple CDS servers would not even be possible (col. 5, lines 29-31) – further evidencing that Claim 3 has been erroneously rejected.

It is further noted that the second alleged server is a redundant server that performs the same function as the primary server, and thus there would be no reason to request information from such a redundant server that the primary server lacks, as is provided by the features of Claim 3 – since it is a redundant server and thus has the same information as the other server such that it can assume the ‘master’ role in the event of a failure (col. 7, lines 1-7).

Further with respect to Claim 3, such claim recites, “receiving, **at a distributed service manager**, a request for a networked service **from a local service manager** for which the local service manager lacks information”. As can be seen, this aspect of Claim 3 is directed to interaction between a distributed service manager and a local service manager, where **a network service request is received** at a distributed service manager, with such request being from the local service manager for which the local service manager lacks information.

The Examiner fails to address such ‘receiving’ aspect of Claim 3 in their rejection of such claim. Instead, the Examiner merely alleges that Elnozahy teaches two initializing steps, and Chandra teaches distributed application caching (see, e.g., pages 9-10 of the Final Office Action dated January 6, 2010). This failure to allege or address the claimed ‘**receiving**’ step is likely because of problems associated with the improper double counting with respect to the Elnozahy CDS service provided by the CDS server. Because the Examiner equates Elnozahy’s CDS as being equivalent to the claimed ‘local service agent’ and Elnozahy’s CDS server as being equivalent to the claimed ‘distributed service manager’, in order for Elnozahy to teach the claimed receiving step, it would have to describe receiving, at the CDS server, a request for a network service from the CDS for which the CDS lacks information. Because the CDS server provides the CDS service, the CDS service would never initiate a request to its own CDS server for information *that it lacks* because if it lacks the information, so would the CDS server as the CDS server provides the functionality of the CDS service – restated, it would not ask itself for

information it does not have since it does not have the information. Thus, it is further urged that Claim 3 has been erroneously rejected due to this additional prima facie obviousness deficiency – as the Examiner fails to even allege, must less establish a teaching of – “receiving, **at a distributed service manager**, a request for a networked service **from a local service manager** *for which the local service manager lacks information* (emphasis added by Appellants).

2. Claim 4

Appellants initially urge error in the rejection of Claim 4 for reasons given above with respect to Claim 3 (of which Claim 4 depends upon).

Further with respect to Claim 4, such claim recites, “sending a request for a networked service from a requesting client to a local service manager associated with the requesting client; and returning information for referencing a matching networked service from the local service manager to the requesting client, wherein the matching networked service has characteristics that match parameters in the request for a networked service”. As can be seen, the features of Claim 4 are directed to a requested network service, where a request for such networked service is sent, and information is returned for a *matching networked service*, where the matching network service has characteristics that match parameters in the request.

In rejecting Claim 4, the Examiner cites Chandry col. 6, lines 20-34 as teaching the returning of matching network service information. Appellants show that there, Chandry states:

In step 72, the cache directory receives an incoming query from a user. Using the table stored in the cache directory as described above, the dispatcher in the cache directory may determine if the query is cached in the system in step 74. **If the query has been previously cached, then the cache directory will direct the cached result to the user to fulfill the user's query in step 76.** This is the best solution since there will be minimal latency since the query does not need to be repeated. This solution also saves computing resources since the query does not need to be repeated. If the query has not been cached, then the cache directory determines if the application cache (AC) computer on the same LAN as the user can service the query in step 78. The AC can service the query if it has the appropriate program and data to execute the query.

As can be seen, this cited passage describes the returning of a query result, with the query result that is returned being the result of a query that was previously run against a database and is now in a local cache (Chandry col. 2, lines 15-51). One example given of such database query result

is a list of companies whose stock price is between \$50 and \$55 (Chandry col. 3, lines 54-61). The returning of the result of a database query - as described by Chandry at the cited passage at col. 6, lines 20-34 - does not teach or suggest returning information for a matching network service, as claimed.

Curiously, in the Examiner's rebuttal comments on page 4 of the Final Office Action dated January 6, 2010, the Examiner now takes the position that the cited Elnozahy – and **not** the cited Chandry reference that the Examiner cites in the 'main body' rejection on page 10 of this same Office Action – teaches all features of Claim 4 due to the two CDS servers ('one CDS server' and 'another CDS server'). Specifically, the Examiner states:

“When one CDS server does not possess information about a networked service, it contacts another CDS server to query for the information.”

Appellants urge that not only does the Examiner not provide any citation of where such teaching is allegedly taught, the Elnozahy reference does not in fact describe such an operation – as the two CDS servers *are redundant with respect to one another*, and *one is not aware of the other* (Elnozahy 'Master' and 'Server' CDS servers, as shown in Figure 1 and described at col. 5, lines 29-31; col. 2, lines 20-28; col. 3, lines 37-46; and col. 5, lines 62-63: “agents 68 and 76 *manage replication*”).

This redundant information between the Master CDS server and Redundant CDS server is highly desired in order to overcome deficiencies with past systems. This can be seen by Elnozahy's discussion regarding failover, when a redundant server takes over for a failed master server, such as at col. 7, lines 1-7:

Continuous availability of lookup and update requests is thus provided, because the **new master is immediately ready to accept CDS requests**. Further, **since the new master had an up-to-date version of each directory entry**, clients are guaranteed to obtain coherent data from the new master, contrasting the operation of vanilla CDS where clients may obtain stale or incoherent information in similar situations.

The Examiner also mischaracterizes Appellants' remarks regarding Claim 4. Specifically, Appellants argued that Chandry did not teach/suggest the features of Claim 4 (see, e.g., pages 9-10

of the Response to Office Action dated August 26, 2009), and yet the Examiner characterizes Appellants' remarks as being:

“ 7. Applicants argues that Elnozahy fails to disclose sending a request ...”

Such mischaracterization further evidences the Examiner's misinterpretation of Appellants' arguments that were presented, and the resulting rebuttal response that changes what reference is alleged to teach the features of Claim 4. Thus, it is further shown that even with this bifurcated rejection (where one reference is alleged to teach all features of Claim 4 in the 'main body' rejection, and another reference is alleged to teach all features of Claim 4 in the "Response to Arguments" section of the same Office Action), the Examiner has still failed to establish prima facie obviousness with respect to such claim – strongly evidencing that such claim is non-obvious in view of the cited references.

3. *Claim 7*

Appellants initially traverse the rejection of Claim 7 for reasons given above with respect to Claim 3 (of which Claim 7 depends upon).

Further with respect to Claim 7, such claim recites, “responsive to a determination that the distributed service manager does not have information about one or more matching networked services, broadcasting the request for a networked service from the distributed service manager to all distributed service managers in the distributed data processing system”. As can be seen, in response to it being determined that the distributed service manager does not have information about a matching service, the request for a networked service *is broadcast to all distributed managers* in the distributed data processing system.

In rejecting Claim 7, the Examiner alleges that all such broadcasting aspects of such claim are taught by Chandra at col. 6, lines 35-61. Appellants show that there, Chandra states:

If the AC local to the user has the appropriate facilities, then the AC executes the query in step 80 and returns the query results to the user. The local AC may also store the query results and forward the query results to the cache directory so that similar future queries may be more easily satisfied. **If the AC local to the user cannot server the query, then the cache directory may select an AC near the user and download/forward the necessary program and data to execute the query in step 82**

so that that AC can execute the query. This option is the least desirable since it requires forwarding the program and data to the AC, but ensures that the query is executed as close to the user as possible so that the latency due to distance between the user and the AC is minimized. In step 84, the cache directory updates its table to reflect that the particular AC has the necessary programs and data to execute the query. In step 86, the AC near the user executes the query. In step 88, the AC forwards the results to the user and stores a copy of the results on the AC. The AC may also forward the results to the cache directory. In this manner, each service request received by the application caching system is executed in the most efficient manner without requiring a copy of the program and data at each node in the network. Thus, the application caching system reduces the amount of redundancy in the network, but ensures that all service requests are handled in the most efficient manner. Now, an example of the application caching in accordance with the invention will be described.

As can be seen, this cited passage describes a conditional step being performed (download/forward the necessary program and data to another AC so that another AC can execute a query), with the triggering condition being that the AC local to the user cannot server (sic – should be ‘service’) the query. It is urged that both the triggering condition and the resulting action are different from what is recited in Claim 7.

For example, the Chandra ‘condition’ is whether the AC local cannot itself *service the query*. In contrast, the claimed ‘condition’ is whether the distributed service manager does not have *information about one or more of the matching network services*. As another example, the Chandra ‘resulting action’ is the downloading/forwarding of necessary programs and data to another computer such that the other computer can execute the query. In contrast, the claimed ‘resulting action’ is ***broadcasting the request*** for a networked service from the distributed service manager to *all* distributed service managers in the distributed data processing system. Thus, it is further urged that Claim 7 has been erroneously rejected due to this additional prima facie obviousness deficiency.

Curiously, in the Examiner’s rebuttal comments beginning on the bottom of page 4 of the Final Office Action dated January 6, 2010, the Examiner now takes the position that the cited Elnozahy – and **not** the cited Chandry reference that the Examiner cites in the ‘main body’ rejection on page 11 of this same Office Action – teaches all features of Claim 7 at col. 7, lines 8-9. Appellants urge that this newly cited passage describes that a new agent broadcasts a ‘designation change’ to all ‘client agents’ on the local area network. This is different from what is recited in Claim 7 for several reasons. First, this broadcasting is not done “responsive to a determination that

the distributed service manager *does not have information about one or more matching networked services*”, as claimed. Instead, such broadcasting occurs when a master server has failed and a redundant server has taken over as the master (Elnozahy col. 6, line 54 – col. 7, line 9). Second, ‘what’ is broadcast is very different. Per Claim 7, it is ‘the request for a network service’ that is broadcast, whereas this newly cited Elnozahy passage broadcasts the master server ‘designation change’. Third, the ‘things’/‘elements’ that the request is broadcast to is different. Per Claim 7, the request is broadcast to ‘all distributed service managers in the distributed data processing system’, whereas this newly cited Elnozahy passage broadcasts to all ‘client agents’ on the local network – which are not CDS servers (which the Examiner alleges is equivalent to the claimed ‘distributed service managers’) but instead are *the requesting clients that make requests to such CDS servers* (Elnozahy Figure 1, element 78; col. 5, lines 7-10). Thus, even this newly cited passage in the Elnozahy reference that is now alleged as teaching all aspects of Claim 7 is also deficient in establishing prima facie obviousness.

The Examiner also mischaracterizes Appellants’ remarks regarding Claim 7. Specifically, Appellants argued that Chandry did not teach/suggest the features of Claim 7 (see, e.g., pages 11-12 of the Response to Office Action dated August 26, 2009), and yet the Examiner characterizes Appellants’ remarks as being:

“ 8. Applicants argues that Elnozahy fails to disclose responsive to a determination ...”

Such mischaracterization further evidences the Examiner’s misinterpretation of Appellants’ arguments that were presented, and the resulting rebuttal response that changes what reference is alleged to teach the features of Claim 7. Thus, it is further shown that even with this bifurcated rejection (where one reference is alleged to teach all features of Claim 7 in the ‘main body’ rejection, and another reference is alleged to teach all features of Claim 7 in the “Response to Arguments” section of the same Office Action), the Examiner has still failed to establish prima facie obviousness with respect to such claim – strongly evidencing that such claim is non-obvious in view of the cited references.

4. *Claim 38*

Appellants initially urge error in the rejection of Claim 38 for reasons given above with respect to Claim 3 (of which Claim 38 depends upon).

Further with respect to Claim 38, such claim recites, “determining, based on the request, whether to return a single matched network service of the set of matched network services or the set of matched network services”. As can be seen, a determination is made *whether to return a single matched network service or a set of matched network services*, where such determination is based on the *request for the network service*.

In rejecting Claim 38, the Examiner alleges that Chandra teaches such single/set determination at col. 6, lines 20-34. Appellants show that there, Chandra states:

In step 72, the cache directory receives an incoming query from a user. Using the table stored in the cache directory as described above, the dispatcher in the cache directory may determine if the query is cached in the system in step 74. If the query has been previously cached, then the cache directory will direct the cached result to the user to fulfill the user's query in step 76. This is the best solution since there will be minimal latency since the query does not need to be repeated. This solution also saves computing resources since the query does not need to be repeated. If the query has not been cached, then the cache directory determines if the application cache (AC) computer on the same LAN as the user can service the query in step 78. The AC can service the query if it has the appropriate program and data to execute the query.

As can be seen, this cited passage describes an incoming ‘query’, where a determination is made as to whether the query is already cached. If so, the cached ‘result’ of the query is returned to the user. If not, a determination is made if the AC computer can service the query. It is urged that a teaching of (1) returning a cached query result, if cached, and (2) determining whether another computer can service a query, if not cached does not teach or suggest a determination is made whether to return a *single matched network service* or a *set of matched network services*, where such determination is based on the request for the network service. Thus, it is further urged that Claim 38 has been erroneously rejected due to this additional prima facie obviousness deficiency.

In the rebuttal section on page 5, the Examiner again attempts to re-introduce new reasons and new references (Elnozahy) in such rebuttal comments when the ‘main body’ alleges the Chandra is alleged to teach the features of Claim 38. In the interest of brevity, Appellants will not repeat the errors in such bifurcated rationale, but instead rely upon similar reasons to those

given above with respect to Claims 4 and 7, and associated mischaracterizations of Appellants remarks that were instead directed toward the teachings of Chandra, and not Elnozahy. It is also noteworthy that the Examiner has failed to provide any specific citation as to where these new alleged Elnozahy coherency cycle teachings are described – further evidencing prima facie obviousness deficiencies. This new allegation also fails to establish a teaching or suggestion that a determination is made as to what to return is ‘based on the request’ for a networked service. Instead, such allegation merely alleges that what is returned is due to a ‘failure’ of a CDS server.

5. Claim 39

Appellants traverse the rejection of Claim 39 for reasons given above with respect to Claim 3 (of which Claim 39 depends upon).

Further with respect to Claim 39, such claim recites, “wherein a plurality of types of networked services are available in the distributed data processing system, and wherein one of the characteristics of a matching service is a type of service”. As can be seen, a plurality of ‘types’ of networked services are available, and one of the characteristics of a matching service is a ‘type of service’.

The Examiner alleges that these characteristics pertaining to a matching service – that is a ‘type’ of service – is described by Chandra at col. 6, lines 20-34. Appellants show that there, Chandra states:

In step 72, the cache directory receives an incoming query from a user. Using the table stored in the cache directory as described above, the dispatcher in the cache directory may **determine if the query is cached** in the system in step 74. If the query has been previously cached, then the cache directory will direct the cached result to the user to fulfill the user's query in step 76. This is the best solution since there will be minimal latency since the query does not need to be repeated. This solution also saves computing resources since the query does not need to be repeated. If the query has not been cached, then the cache directory determines if the application cache (AC) computer on the same LAN as the user can service the query in step 78. The AC can service the query if it has the appropriate program and data to execute the query.

As can be seen, this cited passage describes a cache determination being made in response to receiving a request, such that the received query request does not have to be re-executed if the result to such query is cached. This cited passage does not describe different types of services, or that one of the characteristics of a matching service is a type of service. Thus, it is further urged

that Claim 39 has been erroneously rejected due to this additional prima facie obviousness deficiency.

C. GROUND OF REJECTION 3 (Claims 8-12 and 41)

Claims 8-12 and 41 stand rejected under 35 U.S.C. § 103 as being unpatentable over Elnozahy in view of Chandra, as applied to Claim 8, and further in view of Jindal et al. (U.S. Patent No. 6,324,580), hereinafter “Jindal”.

1. Claims 8-10

Appellants urge error in the rejection of Claim 8 (and similarly for Claims 9 and 10) for similar reasons to those given above with respect to Claim 3 (of which Claim 8 depends upon).

2. Claims 11 and 12

Appellants initially urge error in the rejection of Claim 11 (and similarly for Claim 12) for reasons given above with respect to Claim 10 (of which Claim 11 depends upon).

Further with respect to Claim 11 (and similarly for Claim 12), such claim recites, “comparing one or more of network-related metrics associated with an entire network path between a requesting client and a providing server”. As can be seen, the features of Claim 11 pertain to *network-related metrics* associated with an entire network *path between a requesting client and a providing server* that is used for the comparing step of Claim 10.

In rejecting Claim 11, the Examiner cites Jindal col. 6, lines 8-15 as teaching all aspects of Claim 11. Appellants show that there, Jindal states:

The various pieces of information that may be collected illustratively include: whether a server or instance of a replicated service is operational; the response time for a request submitted to a server or service instance; the number of requests processed by or pending on a server or service instance, a server's proximity (e.g., the number of network hops necessary to reach the server from nameserver 100), etc.

As can be seen, this cited passage describes the various pieces of information that may be collected, including: (1) whether a server is operational; (2) the response time for a request submitted to a server; (3) the number of server requests; and (4) a server's proximity to the nameserver. However,

all of the characteristics are specific to the server itself. None of these described pieces of information are described as being associated with a particular network path between a requesting client (depicted by element 120 of Jindal Figure 1) and a server (depicted by elements 110, 112 and 114 of Figure 1). At best, this cited passage describes use of the network path between the nameserver 100 of Jindal Figure 1 and the servers 110, 112 and 114 of Figure 1. Because Jindal does not describe any ability to obtain network characteristics pertaining to network 122 of Figure 1 - of which a requesting client attaches to - it would have no way of implementing the client/server path characteristics expressly recited in Claim 11. Thus, it is further urged that Claim 11 has been erroneously rejected due to this additional prima facie obviousness deficiency.

3. Claim 41

Appellants initially urge error in the rejection of Claim 41 for reasons given above with respect to Claim 7 (of which Claim 41 depends upon).

Further with respect to Claim 41, such claim recites “wherein each of the distributed service managers includes a localization module, wherein the parameters within respective localization modules are tailored to provide different load balancing for corresponding distributed service managers”. As can be seen, the features of Claim 41 are directed to particular characteristics of the distributed service managers, and in particular that *each of the distributed service managers includes a localization module having load balancing parameters*.

In rejecting Claim 41, the Examiner alleges that all features of Claim 41 are taught by Jindal at column 4, lines 49-67. Appellants show that there, Jindal states:

In yet another embodiment of the invention, status objects 200a, 200b and 200c are configured to determine whether a particular service (e.g., web service, electronic mail service), application program or server is operational. Illustratively, the status objects in this embodiment issue a Connect (or similar) command to the target service or server. If a Connect command is successful the issuing object knows that the target is operational, otherwise it is assumed to be inoperative.

Illustratively, for each replicated service (or application) that is to be monitored (i.e., that is subject to load balancing) on a server, a separate status object operates on nameserver 100. In addition, each status object illustratively performs a single function (e.g., determine response time, determine a server's distance from nameserver 100). In alternative embodiments of the invention, however, a single status object may monitor multiple servers or services and/or perform multiple functions.

In FIG. 2, individual monitor objects (IMO) 202a, 202b and 202c also reside and execute on nameserver 100. A separate IMO is depicted for each instance of the replicated service.

This cited passage describes a plurality of status objects that each performs a single function. These status objects are not described as having a localization module having parameters therein that are tailored to provide different load balancing *for corresponding distributed service managers*. Instead, they provide a single status acquisition function (col. 6, lines 59-60).

This cited passage also describes a plurality of individual monitor objects. However, these monitor objects are not described as having a localization module having parameters therein that are tailored to provide different load balancing *for corresponding distributed service managers*. Instead, they provide status object invocation and the collecting of information from the status objects (col. 7, lines 1-4).

Thus, this cited passage does not describe wherein *each* of the distributed service managers *includes a localization module*, wherein the parameters within respective localization modules are tailored to provide different load balancing for corresponding distributed service managers, as specifically recited in Claim 41. Thus, it is further urged that Claim 41 has been erroneously rejected due to this additional prima facie obviousness deficiency.

D. GROUND OF REJECTION 4 (Claim 37)

Claim 37 stands rejected under 35 U.S.C. § 103 as being unpatentable over Elnozahy in view of Chandra, as applied to Claim 3, and further in view of Fowlow et al. (U.S. Patent No. 5,920,868), hereinafter “Fowlow”.

With respect to Claim 37, such claim recites, “configuring the local service manager to not provide access to object request broker (ORB) services that provide internal service and which are valid only in a scope of a local ORB; configuring the local service manager to provide access to ORB services that are instantiated on each ORB only through requests based on an ORB identifier; and configuring the local service manager to provide access to ORB services that may be accessed from outside the scope of the local ORB through requests based on both a service specification string and an ORB identifier”. As can be seen, the features of Claim 37 are directed to the *interplay between* the ‘local service manager’ and the ‘object request broker’.

In rejecting Claim 37, the Examiner states that Elnozahy in view of Chandra fail to disclose the use of object request broker services in accessing services in a distributed environment. The Examiner goes on to state the Fowlow discloses ‘accessing objects in a distributed system using an ORB service’. However, even assuming such assertion is true, such generalized assertion does not establish a teaching or suggestion pertaining to the specific features recited in Claim 37, such as *specially configuring the local service manager* to either (i) provide access to ORB services or (ii) not provide access to ORB services *based upon particular ORB characteristics*, including (1) object request broker (ORB) services that provide internal service and which are valid only in a scope of a local ORB, (2) ORB services that are instantiated on each ORB only through requests based on an ORB identifier, and (3) ORB services that may be accessed from outside the scope of the local ORB through requests based on both a service specification string and an ORB identifier - as is specifically recited in Claim 37. Thus, Claim 37 has been erroneously rejected as the Examiner has failed to properly establish prima facie obviousness.

Appellants further traverse the rejection of Claim 37 for similar reasons to those given above with respect to Claim 3, as such claim includes the features of Claim 3.

E. CONCLUSION

As shown above, the Examiner has failed to state valid rejections against any of the claims. Therefore, Appellants request that the Board of Patent Appeals and Interferences reverse the rejections. Additionally, Appellants request that the Board direct the examiner to allow the claims.

Date: May 4, 2010

Respectfully Submitted,

/Wayne P. Bailey/

Wayne P. Bailey
Reg. No.34,289
Yee & Associates, P.C.
P.O. Box 802333
Dallas, TX 75380
(972) 385-8777

CLAIMS APPENDIX

The text of the claims involved in the appeal is as follows:

3. A method of balancing demand for networked services in a distributed data processing system, the method comprising the steps of:

initializing one or more local service managers within the distributed data processing system, wherein each local service manager has information about and provides access to networked services defined within a respective local region of the distributed data processing system for clients within the distributed data processing system, and wherein each client is uniquely associated with a local service manager;

initializing one or more distributed service managers within the distributed data processing system, wherein each distributed service manager provides access to the networked services to the local service managers within the distributed data processing system, and wherein each local service manager is uniquely associated with a distributed service manager;

receiving, at a distributed service manager, a request for a networked service from a local service manager for which the local service manager lacks information;

determining whether the distributed service manager has information about a networked service with one or more characteristics that match one or more parameters in the request for a networked service, wherein the determining step is accomplished by reference to a cache maintained by the distributed service manager which contains information resulting from prior requests for networked services; and

returning information for referencing a matched networked service.

4. The method of claim 3 further comprising:
 - sending a request for a networked service from a requesting client to a local service manager associated with the requesting client; and
 - returning information for referencing a matching networked service from the local service manager to the requesting client, wherein the matching networked service has characteristics that match parameters in the request for a networked service.
5. The method of claim 3 further comprising:
 - receiving a request for a networked service at a local service manager; and
 - determining whether the local service manager has information for referencing a networked service with characteristics that match parameters in the request for a networked service.
6. The method of claim 5 further comprising:
 - responsive to a determination that the local service manager has information about a matching networked service, returning the information for referencing the matching networked service to the requesting client;
 - responsive to a determination that the local service manager does not have information about a matching networked service, forwarding the request for a networked service from the local service manager to a distributed service manager associated with the local service manager.

7. The method of claim 3 further comprising:

responsive to a determination that the distributed service manager does not have information about one or more matching networked services, broadcasting the request for a networked service from the distributed service manager to all distributed service managers in the distributed data processing system;

receiving information for referencing one or more matching networked services at the distributed service manager in response to the broadcast request; and

caching the received information for referencing one or more matching networked services at the distributed service manager.

8. The method of claim 3 further comprising:

in response to a determination that the distributed service manager has information about two or more matching networked services, selecting a single networked service at the distributed service manager.

9. The method of claim 8 further comprising:

performing a load balancing operation at the distributed service manager to select the single networked service.

10. The method of claim 9 further comprising:

comparing network-related metrics during the load balancing operation.

11. The method of claim 10 further comprising:
comparing one or more of network-related metrics associated with an entire network path between a requesting client and a providing server.
12. The method of claim 11 wherein the network-related metrics are realtime network-related metrics and are selected from a group comprising:
bottleneck-link speed, round-trip time, and hop count.
37. A method of balancing demand for networked services in a distributed data processing system, the method comprising the steps of:
initializing one or more local service managers within the distributed data processing system, wherein each local service manager has information about and provides access to networked services defined within a respective local region of the distributed data processing system for clients within the distributed data processing system, and wherein each client is uniquely associated with a local service manager;
initializing one or more distributed service managers within the distributed data processing system, wherein each distributed service manager provides access to the networked services to the local service managers within the distributed data processing system, and wherein each local service manager is uniquely associated with a distributed service manager;
receiving, at a distributed service manager, a request for a networked service from a local service manager for which the local service manager lacks information;
determining whether the distributed service manager has information about a networked

service with one or more characteristics that match one or more parameters in the request for a networked service, wherein the determining step is accomplished by reference to a cache maintained by the distributed service manager which contains information resulting from prior requests for networked services;

returning information for referencing a matched networked service;

configuring the local service manager to not provide access to object request broker (ORB) services that provide internal service and which are valid only in a scope of a local ORB;

configuring the local service manager to provide access to ORB services that are instantiated on each ORB only through requests based on an ORB identifier; and

configuring the local service manager to provide access to ORB services that may be accessed from outside the scope of the local ORB through requests based on both a service specification string and an ORB identifier.

38. The method of claim 3 further comprising:

determining whether the distributed service manager has information about a plurality of networked services with characteristics that match parameters in the request for a networked service and forming a set of matched network services;

determining, based on the request, whether to return a single matched network service of the set of matched network services or the set of matched network services;

responsive to a determination to return a single matched network service, returning information for referencing the single matched networked service from the distributed service manager to the local service manager; and

responsive to a determination to return the set of matched network services, returning

information for referencing the set of matched network services from the distributed service manager to the local service manager.

39. The method of claim 3 wherein a plurality of types of networked services are available in the distributed data processing system, and wherein one of the characteristics of a matching service is a type of service.

40. The method of claim 7 wherein each of the distributed service managers caches information resulting from requests of supported clients, and wherein the information which respective distributed service manager differs according to the requests of supported clients.

41. The method of claim 7 wherein each of the distributed service managers includes a localization module, wherein the parameters within respective localization modules are tailored to provide different load balancing for corresponding distributed service managers.

EVIDENCE APPENDIX

This appeal brief presents no additional evidence.

RELATED PROCEEDINGS APPENDIX

This appeal has no related proceedings.